

EAGLE can do Efficient LTL Monitoring

Howard Barringer^{*1}, Allen Goldberg², Klaus Havelund² and Koushik Sen^{**3}

¹ University of Manchester, England

² Kestrel Technology, NASA Ames Research Center, USA

³ University of Illinois, Urbana Champaign, USA

Abstract. We briefly present a rule-based framework, called EAGLE, that has been shown to be capable of defining and implementing finite trace monitoring logics, including future and past time temporal logic, extended regular expressions, real-time logics, interval logics, forms of quantified temporal logics, and so on. In this paper we show how EAGLE can do linear temporal logic (LTL) monitoring in an efficient way. We give an upper bound on the space and time complexity of this monitoring.

1 Introduction

Runtime verification, or runtime monitoring, comprises having a software module, an observer, monitor the execution of a program, and check its conformity with a requirement specification, often written in a temporal logic or as a state machine. Runtime verification can be applied to evaluate automatically test runs, either on-line or off-line, analyzing stored execution traces; or it can be used on-line during operation, potentially steering the application back to a safety region if a property is violated. It is highly scalable. Several runtime verification systems have been developed, of which some were presented at three recent international workshops on runtime verification [1].

Linear temporal logic (LTL) [17] has been core to several of these attempts. The commercial tool Temporal Rover (TR) [5, 6] supports a fixed future and past time LTL, with the possibility of specifying real-time and data constraints (time-series) as annotations on the temporal operators. Its implementation is based on alternating automata. Algorithms using alternating automata to monitor LTL properties are also proposed in [8], and a specialized LTL collecting statistics along the execution trace is described in [7]. The MAC logic [16] is a form of past-time LTL with operators inspired by interval logics and which models real-time via explicit clock variables. A logic based on extended regular expressions [18] has also been proposed and is argued to be more succinct for certain properties. The logic described in [14] is a sophisticated interval logic, argued to be more user-friendly than plain LTL. Our own previous work includes the development of several algorithms, such as generating dynamic programming algorithms for past time logic [12], using a rewriting system for monitoring future-time logic [11], or generating Büchi automata inspired algorithms adapted to finite trace LTL [10].

* This author is most grateful to RIACS/USRA and to the UK's EPSRC under grant GR/S40435/01 for the partial support provided to conduct this research.

** This author is grateful for the support received from RIACS to undertake this research while participating in the Summer Student Research Program at the NASA Ames Research Center.

This large variety of logics prompted us to search for a small and general framework for defining monitoring logics, which would be powerful enough to capture essentially all of the above described logics, hence supporting future and past time logics, interval logics, extended regular expressions, state machines, real-time and data constraints, and statistics. The framework should support the definition of new logics in an easy manner and should support the monitoring of programs with their complex program states. The result of our search is the logic EAGLE which is described in details in [4]. In this paper we briefly describe EAGLE and its expressivity and then focus mainly on the subset LTL and analyze its complexity. The EAGLE logic and its implementation for run-time monitoring has in particular been significantly influenced by earlier work of Barringer et al., see for example [3], on the executable temporal logic METATEM. A linear-time temporal formula can be separated [9] into a boolean combination of pure past, present and pure future time formulas - in particular, the combination can be written as a collection of "directly executable" global conditional rules of the form "if pure past-time then present-time and pure future-time". The present-time, or state, formulas determine how the state for the current moment in time is built and the pure future-time formulas yield obligations that need to be fulfilled at some time later. The separation result, rules and future obligations are central in our current work. However, the fundamental difference between METATEM and EAGLE is that the METATEM interpreter builds traces state by state, whereas EAGLE is used for checking given finite traces: costly implementation features, such as backtracking and loop-checking, are not required.

We recently discovered parallel work [15] using recursive equations to implement a real-time logic. However we had already developed the ideas further. We provide the language of recursive equations to the user, we support a mixture of future time and past time operators, we treat real-time as a special case of data values, and hence we allow a very general logic for reasoning about data, including the possibility of relating data values across the execution trace, both forwards and backwards.

The paper is structured as follows. Section 2 introduces our logic framework, then in section 3 we discuss the algorithm and calculus that underlies our implementation for the special case of LTL, which is then briefly described along with complexity and initial experimentation in section 4.

2 The Logic

In this section we briefly describe the temporal finite trace monitoring logic EAGLE[4]. The logic offers a succinct but powerful set of primitives, essentially supporting recursive parameterized equations, with a minimal/maximal fix-point semantics together with three temporal operators: next-time, previous-time, and concatenation. The next-time and previous-time operators can be used for defining future time respectively past time temporal logics on top of EAGLE. The concatenation operator can be used to define interval logics and an extended regular expression language. Rules are parameterized to allow for reasoning about data values, including real-time. Atomic propositions are boolean expressions over a program state, Java states in the current implementation. The logic is first introduced informally through two examples whereafter its syntax and semantics is given. Finally, its relationship to some other important logics is outlined.

2.1 EAGLE by Example

Fundamental Concepts Assume we want to state a property about a program P , which contains the declaration of two integer variables x and y . We want to state that whenever x is positive then eventually y becomes positive. The property can be written as follows in classical future time LTL: $\Box(x > 0 \rightarrow \Diamond y > 0)$. The formulas $\Box F$ (always F) and $\Diamond F$ (eventually F), for some property F , usually satisfy the following equivalences, where the temporal operator $\bigcirc F$ stands for *next* F (meaning 'in next state F ')

$$\Box F \equiv F \wedge \bigcirc(\Box F) \quad \Diamond F \equiv F \vee \bigcirc(\Diamond F)$$

One can show that $\Box F$ is a solution to the recursive equation $X = F \wedge \bigcirc X$; in fact it is the maximal solution. A fundamental idea in our logic is to support this kind of recursive definition, and to enable users define their own temporal combinators using equations similar to those above. In the current framework one can write the following definitions for the two combinators Always and Eventually, and the formula to be monitored (M_1):

$$\begin{aligned} \underline{\max} \text{ Always}(\text{Form } F) &= F \wedge \bigcirc \text{Always}(F) \\ \underline{\min} \text{ Eventually}(\text{Form } F) &= F \vee \bigcirc \text{Eventually}(F) \\ \underline{\text{mon}} M_1 &= \text{Always}(x > 0 \rightarrow \text{Eventually}(y > 0)) \end{aligned}$$

The Always operator is defined as a maximal fix-point operator; the Eventually operator is defined as a minimal fix-point operator. Maximal rules define safety properties (nothing bad ever happens), while minimal rules define liveness properties (something good eventually happens). For us, the difference only becomes important when evaluating formulas at the boundaries of a trace. To understand how this works it suffices to say here that monitored rules evolve as new states are appearing. Assume that the end of the trace has been reached (we are beyond the last state) and a monitored formula F has evolved to F' . Then all applications in F' of maximal fix-point rules will evaluate to true, since they represent safety properties that apparently have been satisfied throughout the trace, while applications of minimal fix-point rules will evaluate to false, indicating that some event did not happen. Assume for example that we evaluate the formula M_1 in a state where $x > 0$ and $y \leq 0$, then as a liveness obligation for the future we will have the expression:

$$\bigcirc \text{Eventually}(y > 0) \wedge \bigcirc \text{Always}(x > 0 \rightarrow \text{Eventually}(y > 0))$$

Assume that we at this point detect the end of the trace; that is: we are beyond the last state. The outstanding liveness obligation $\text{Eventually}(y > 0)$ has not yet been fulfilled, which is an error. This is captured by the evaluation of the minimal fix-point combinator Eventually to false at this point. The remaining other obligation from the \wedge -formula, namely, $\text{Always}(x > 0 \rightarrow \text{Eventually}(y > 0))$, is a safety property and evaluates to true.

For completeness we provide remaining definitions of the future time LTL operators \mathcal{U} (until) and \mathcal{W} (unless) below. Note how \mathcal{W} is defined in terms of other operators. However, it could have been defined recursively.

$$\begin{aligned} \underline{\min} \text{ Until}(\text{Form } F_1, \text{Form } F_2) &= F_2 \vee (F_1 \wedge \bigcirc \text{Until}(F_1, F_2)) \\ \underline{\max} \text{ Unless}(\text{Form } F_1, \text{Form } F_2) &= \text{Until}(F_1, F_2) \vee \text{Always}(F_1) \end{aligned}$$

Data Parameters We have seen how rules can be parameterized with formulas. Let's complicate the example with data parameters. Suppose we want to state the property: "whenever at some point $k = x > 0$ for some k , then eventually $y == k$ ". This can be stated as follows in quantified LTL: $\Box(x > 0 \rightarrow \exists k.(k = x \wedge \Diamond y = k))$. We use parameterized rules to state this property, capturing the value of x when $x > 0$ as a rule parameter.

$$\underline{\min} R(\underline{\text{int}} k) = \text{Eventually}(s.y == k) \quad \underline{\text{mon}} M_2 = \text{Always}(s.x > 0 \rightarrow R(s.x))$$

Rule R is parameterized with an integer k , and is instantiated in M_2 when $x > 0$, hence capturing the value of x at that moment. Rule R replaces the existential quantifier. The logic also provides a previous-time operator, which allows us to define past time operators; the data parametrization works uniformly for rules over past as well as future, which is non-trivial to achieve since the implementation does not store the trace, see Section 4. Data parametrization is also used to elegantly model real-time logics.

2.2 Syntax and Semantics

Syntax A specification S consists of a declaration part D and an observer part O . D consists of zero or more rule definitions R , and O consists of zero or more monitor definitions M , which specify what to be monitored. Rules and monitors are named (N).

$$\begin{aligned} S^c &::= \text{dec } D \text{ obs } O \\ D &::= R^* \\ O &::= M^* \\ R &::= \{\underline{\max} \mid \underline{\min}\} N(T_1 x_1, \dots, T_n x_n) = F \\ M &::= N = F \\ T &::= \underline{\text{Form}} \mid \text{java primitive type} \\ F &::= \text{java expression} \mid \underline{\text{true}} \mid \underline{\text{false}} \mid \neg F \mid F_1 \wedge F_2 \mid F_1 \vee F_2 \mid F_1 \rightarrow F_2 \mid \\ &\quad \bigcirc F \mid \odot F \mid F_1 \cdot F_2 \mid N(F_1, \dots, F_n) \end{aligned}$$

A rule definition R is preceded by a keyword indicating whether the interpretation is maximal or minimal (which we recall determines the value of a rule application at the boundaries of the trace). Parameters are typed, and can either be a formula of type Form, or of a primitive Java type, such as int, long, float, etc.. The body of a rule/monitor is a formula of the syntactic category *Form* (with meta-variables F , etc.). The propositions of this logic are Java expressions over an observer state. These can be arbitrary Java expressions using all of Java's expression language constructs, recommended not to have no side effects. Formulas are composed using standard propositional logic operators together with a next-state operator ($\bigcirc F$), a previous-state operator ($\odot F$), and a concatenation-operator ($F_1 \cdot F_2$). Finally, rules can be applied and their parameters must be type correct; formula arguments can be any formula, with the exception that if an argument is a java expression, it must be of boolean type.

Semantics The semantics of the logic is defined in terms of a satisfaction relation $\models \subseteq \text{Trace} \times \text{Form}$ between execution traces and specifications. An execution trace σ

is a finite sequence of program states $\sigma = s_1 s_2 \dots s_n$, where $|\sigma| = n$ is the length of the trace. The i 'th state s_i of a trace σ is denoted by $\sigma(i)$. The term $\sigma^{[i,j]}$ denotes the subtrace of σ from position i to position j , both positions included. In the implementation a state is a user defined Java object that is updated through a user provided *update* method for each new event generated by the program. Given a trace σ and a specification $\text{dec } D$ $\text{obs } O$, satisfaction is defined as follows:

$$\sigma \models \text{dec } D \text{ obs } O \text{ iff } \forall (N = F) \in O. \sigma, 1 \models_D F$$

That is, a trace satisfies a specification if the trace, observed from position 1 (the first state), satisfies each monitored formula. The definition of the satisfaction relation $\models_D \subseteq (\text{Trace} \times \text{nat}) \times \text{Form}$, for a set of rule definitions D , is presented below, where $0 \leq i \leq n+1$ for some trace $\sigma = s_1 s_2 \dots s_n$. Note that the position of a trace can become 0 (before the first state) when going backwards, and can become $n+1$ (after the last state) when going forwards, both cases causing rule applications to evaluate to either true if maximal or false if minimal, without considering the body of the rules at that point.

$$\begin{array}{ll} \sigma, i \models_D \text{jexp} & \text{iff } 1 \leq i \leq |\sigma| \text{ and } \text{evaluate}(\text{jexp})(\sigma(i)) == \text{true} \\ \sigma, i \models_D \underline{\text{true}} & \\ \sigma, i \not\models_D \underline{\text{false}} & \\ \sigma, i \models_D \neg F & \text{iff } \sigma, i \not\models_D F \\ \sigma, i \models_D F_1 \wedge F_2 & \text{iff } \sigma, i \models_D F_1 \text{ and } \sigma, i \models_D F_2 \\ \sigma, i \models_D F_1 \vee F_2 & \text{iff } \sigma, i \models_D F_1 \text{ or } \sigma, i \models_D F_2 \\ \sigma, i \models_D F_1 \rightarrow F_2 & \text{iff } \sigma, i \models_D F_1 \text{ implies } \sigma, i \models_D F_2 \\ \sigma, i \models_D \bigcirc F & \text{iff } i \leq |\sigma| \text{ and } \sigma, i+1 \models_D F \\ \sigma, i \models_D \odot F & \text{iff } 1 \leq i \text{ and } \sigma, i-1 \models_D F \\ \sigma, i \models_D F_1 \cdot F_2 & \text{iff } \exists j \geq i \text{ s.t. } \sigma^{[1, j-1]}, i \models_D F_1 \text{ and } \sigma^{[j, |\sigma|]}, 1 \models_D F_2 \\ \sigma, i \models_D N(F_1, \dots, F_m) & \text{iff } \begin{cases} \text{if } 1 \leq i \leq |\sigma| \text{ then:} \\ \quad \sigma, i \models_D F[x_1 \mapsto F_1, \dots, x_n \mapsto F_n] \\ \quad \text{where } (N(T_1 x_1, \dots, T_n x_n) = F) \in D \\ \text{otherwise, if } i = 0 \text{ or } i = |\sigma| + 1 \text{ then:} \\ \quad \text{rule } N \text{ is defined as } \underline{\text{max}} \text{ in } D \end{cases} \end{array}$$

A Java expression (a proposition) is evaluated in the current state in case the position i is within the trace ($1 \leq i \leq n$). In the boundary cases ($i = 0$ and $i = n+1$) Java expressions evaluate to false. Propositional operators have their standard semantics in all positions. A next-time formula $\bigcirc F$ evaluates to true if the current position is not beyond the last state and F holds in the next position. Dually for the previous-time formula. This means that these formulas always evaluate to false in the boundary positions (0 and $n+1$). The concatenation formula $F_1 \cdot F_2$ is true if the trace σ can be split into two subtraces $\sigma = \sigma_1 \sigma_2$, such that F_1 is true on σ_1 , observed from the current position i , and F_2 is true on σ_2 (ignoring σ_1 , and thereby limiting the scope of past time operators). Applying a rule within the trace (positions $1 \dots n$) consists of replacing the call with the right-hand side of the definition, substituting arguments for formal parameters. At the boundaries (0 and $n+1$) a rule application evaluates to true if and only if it is maximal.

2.3 Linear Temporal Logic in EAGLE

In this section we define the different operators for LTL in the form of rules. For the special case of EAGLE, one writes an LTL formula to be monitored using these operators only. We do not allow to define new operators in the form of rules, as our algorithm is hardwired with these operators and we do not do any synthesis of monitor for this special case, LTL. However, in general EAGLE, one can define new temporal operators as rules and in that case we synthesize the monitors.

Future Time LTL We start with the standard future time linear temporal logic based on the temporal modalities, \bigcirc and \mathcal{U} . We also define the other modalities, such as \square , \diamond and \mathcal{W} , directly as rules in EAGLE. Our embedding relies upon the fact that the formula $\phi \mathcal{U} \psi$ corresponds to the minimal solution to the equation $X = \psi \vee \phi \wedge \bigcirc X$.

$$\begin{aligned} \underline{\max} \text{ Always}(\text{Form } F) &= F \wedge \bigcirc \text{Always}(F) \\ \underline{\min} \text{ Eventually}(\text{Form } F) &= F \vee \bigcirc \text{Eventually}(F) \\ \underline{\min} \text{ Until}(\text{Form } F_1, \text{Form } F_2) &= F_2 \vee (F_1 \wedge \bigcirc \text{Until}(F_1, F_2)) \\ \underline{\max} \text{ Unless}(\text{Form } F_1, \text{Form } F_2) &= F_2 \vee (F_1 \wedge \bigcirc \text{Unless}(F_1, F_2)) \end{aligned}$$

Note that the unless modality is defined as maximal since we require that $\text{Unless}(F_1, F_2)$ evaluates to true on the empty sequence, unlike $\text{Until}(F_1, F_2)$ that must evaluate to false on the empty sequence.

Past Time LTL: A past time linear temporal logic, i.e. one whose temporal modalities only look to the past, could be defined in the mirror way to the future time logic by using the built-in previous modality, \ominus , in place of the future next time modality, \bigcirc . We define an explicit rule, Previous, for the \ominus modality in Past Time LTL. This is done for technical reason (becomes clear later) and we assume that in an LTL formula one uses Previous instead of \ominus . Note that the Zince rule defines the past-time correspondent to the future time unless, or weak until, modality, i.e. it is a weak version of Since.

$$\begin{aligned} \underline{\min} \text{ Previous}(\text{Form } F) &= \ominus F \\ \underline{\max} \text{ AlwaysInPast}(\text{Form } F) &= F \wedge \ominus \text{AlwaysInPast}(F) \\ \underline{\min} \text{ EventuallyInPast}(\text{Form } F) &= F \vee \ominus \text{EventuallyInPast}(F) \\ \underline{\min} \text{ Since}(\text{Form } F_1, \text{Form } F_2) &= F_2 \vee (F_1 \wedge \ominus \text{Since}(F_1, F_2)) \\ \underline{\max} \text{ Zince}(\text{Form } F_1, \text{Form } F_2) &= F_2 \vee (F_1 \wedge \ominus \text{Zince}(F_1, F_2)) \end{aligned}$$

Combined Future and Past Time LTL: By combining the definitions for the future and past time LTLs defined above, we obtain a temporal logic over the future, present and past, in which one can freely intermix the future and past time modalities (to any depth).

2.4 Relationship to Other Logics

Although in this paper we use EAGLE for LTL monitoring, in general EAGLE is expressively rich; indeed, any linear-time temporal logic, whose temporal modalities can be

recursively defined over the next, past or concatenation modalities, can be embedded within it. Furthermore, since in effect we have a limited form of quantification over possibly infinite data sets, and concatenation, we are strictly more expressive than, say, a linear temporal fixed point logic (over next and previous). A formal characterization of the logic is beyond the scope of this paper, however to make the paper self-contained, we demonstrate the logic's utility and expressiveness through some examples. However, if the readers interested on LTL monitoring can skip this subsection.

Combined Future and Past Time LTL with Data Values: We are thus able to express constraints such as if ever the variable x exceeds 0, there was an earlier moment when the variable y was 4 and then remains with that value until it gets increased sometime later, possibly after the moment when x exceeds 0.

$$\underline{\text{mon}} M_2 = \text{Always}(x > 0 \rightarrow \text{EventuallyInPast}(y == 4 \wedge \text{Until}(y == 4, y > 4)))$$

Extended LTL and μTTL : The ability to define temporal modalities recursively provides the ability to define Wolper's ETL or the semantically equivalent fixpoint temporal calculus. Such expressiveness is required to capture regular properties such as temporal formula F is required to be true on every even moment of time:

$$\underline{\text{max}} \text{Even}(\text{Form } F) = F \wedge \bigcirc \bigcirc \text{Even}(F)$$

The μTTL formula $\forall x.p \wedge \bigcirc \bigcirc x \wedge \mu y.q \wedge \bigcirc x \vee \bigcirc y$, where p and q are atomic formulas, would be denoted by the formula, $X()$, where rules X and Y are:

$$\underline{\text{max}} X() = p \wedge \bigcirc \bigcirc X() \wedge Y() \quad \underline{\text{min}} Y() = q \wedge \bigcirc X() \vee \bigcirc Y()$$

Extended Regular Expressions: The language of Extended Regular Expressions (ERE), i.e. adding complementation to regular expressions, has been proposed as a powerful formalism for run-time monitoring. ERE can straightforwardly be embedded within our rule-based system. Given, $E ::= \emptyset | \varepsilon | a | E \cdot E | E + E | E \cap E | \neg E | E^*$, let $\text{Tr}(E)$ denote the ERE E 's corresponding EAGLE formula. For convenience, we define the rule $\underline{\text{max}} \text{Empty}() = \neg \bigcirc \text{true}$ which is true only when evaluated on an empty (suffix) sequence. Tr is inductively defined as follows.

$$\begin{array}{ll} \text{Tr}(\emptyset) & = \underline{\text{false}} & \text{Tr}(\varepsilon) & = \text{Empty}() \\ \text{Tr}(a) & = a \wedge \bigcirc \text{Empty}() & \text{Tr}(E_1 \cdot E_2) & = \text{Tr}(E_1) \cdot \text{Tr}(E_2) \\ \text{Tr}(E_1 + E_2) & = \text{Tr}(E_1) \vee \text{Tr}(E_2) & \text{Tr}(E_1 \cap E_2) & = \text{Tr}(E_1) \wedge \text{Tr}(E_2) \\ \text{Tr}(\neg E) & = \neg \text{Tr}(E) & & \\ \text{Tr}(E^*) & = X() \text{ where } \underline{\text{max}} X() = \text{Empty}() \mid (\text{Tr}(E) \cdot X()) & & \end{array}$$

Real Time as a Special Case of Data Binding: Metric temporal logics, in which temporal modalities are parameterized by some underlying real-time clock(s), can be straightforwardly embedded into our system through rule parameterization. For example, consider the metric temporal modality, $\diamond^{[t_1, t_2]}$ in a system with just one global clock.

An absolute interpretation of $\diamond^{[t_1, t_2]}\phi$ has the formula true if and only if ϕ holds at some time in the future when the real-time clock has value within the interval $[t_1, t_2]$. For our context, we assume that the finite sequence of states being monitored contains a variable *clock* giving the real-time value of the clock for the associated state. The rule

$$\underline{\min} \text{EventAbs}(\underline{\text{Form}} F, \underline{\text{long}} t_1, \underline{\text{long}} t_2) = \text{clock} \leq t_2 \wedge (F \rightarrow t_1 \leq \text{clock}) \wedge (\neg F \rightarrow \bigcirc \text{EventAbs}(F, t_1, t_2))$$

defines the operator $\diamond^{[t_1, t_2]}$ for absolute values of the clock. A relativized version of the modality can then be defined as:

$$\underline{\min} \text{EventRel}(\underline{\text{Form}} F, \underline{\text{long}} t_1, \underline{\text{long}} t_2) = \text{EventAbs}(F, \text{clock} + t_1, \text{clock} + t_2)$$

Counting and Statistical Calculations: In a monitoring context, one may wish to gather statistics on the truth of some property, for example whether a particular state property ϕ holds with at least some probability p over a given sequence, i.e. it doesn't fail with probability greater than $(1 - p)$. Consider the operator $\Box_p \phi$ defined by:

$$\sigma, i \models \Box_p \phi \text{ iff } \exists S \subseteq \{i..|\sigma|\} \text{ s.t. } \frac{|S|}{|\sigma| - i} \geq p \wedge \forall j \in S. \sigma, j \models \phi$$

An encoding within our logic can then be given as:

$$\begin{aligned} \underline{\min} \text{A}(\underline{\text{Form}} \phi, \underline{\text{float}} p, \underline{\text{int}} f, \underline{\text{int}} t) = & (\bigcirc \text{Empty}()) \wedge ((\phi \wedge (1 - \frac{f}{t}) \geq p) \vee (\neg \phi \wedge (1 - \frac{f+1}{t}) \geq p)) \vee \\ & (\neg \text{Empty}()) \wedge ((\phi \rightarrow \bigcirc \text{A}(\phi, p, f, t+1)) \wedge (\neg \phi \rightarrow \bigcirc \text{A}(\phi, p, f+1, t+1))) \\ \underline{\min} \text{AtLeast}(\underline{\text{Form}} \phi, \underline{\text{float}} p) = & \text{A}(\phi, p, 0, 1) \end{aligned}$$

Towards Context Free: Above we showed that EAGLE could encode logics such as ETL, which extend LTL with regular grammars (when restricted to finite traces), or even extended regular expressions. In fact, we can go beyond regularity into the world of context-free languages, necessary, for example, to express properties such as every login is matched by a logout and at no point are there more logouts than logins. Indeed, such a property can be expressed in several ways in EAGLE. Assume we are monitoring a sequence of *login* and *logout* events. We can define a rule $\text{Match}(\underline{\text{Form}} F_1, \underline{\text{Form}} F_2)$ and monitor with $\text{Match}(\text{login}, \text{logout})$ where:

$$\underline{\min} \text{Match}(\underline{\text{Form}} F_1, \underline{\text{Form}} F_2) = F_1 \cdot \text{Match}(F_1, F_2) \cdot F_2 \cdot \text{Match}(F_1, F_2) \vee \text{Empty}()$$

Less elegantly, and which we leave as an exercise, one could use the rule parametrization mechanism to count the numbers of logins and logouts.

3 Algorithm

In this section, we now outline the computation mechanism used to determine whether a given monitoring formula given in LTL holds for some given input sequence of events.

On the observer side a local state is maintained. The *atomic propositions* are specified with respect to the variables in this local state. At every event the observer modifies the local state of the observer based on that event and then evaluates the monitored formulas on that state and generates a new set of monitored formulas. At the end of the trace the value of the monitored formulas are determined. The evaluation of a formula F on a state $s = \sigma(i)$ in a trace σ results in an another formula $eval(F, s)$ with the property that $\sigma, i \models F$ if and only if $\sigma, i + 1 \models eval(F, s)$. The definition of the operator $eval : Form \times State \rightarrow Form$ uses another auxiliary operator $update : Form \times State \rightarrow Form$. The intuition behind using the operator $update$ is to update a formula properly in presence of previous operators. The value of a formula F at the end of a trace is given by $value(F)$. The operator $value : Form \rightarrow \{\underline{true}, \underline{false}\}$ returns true if the formula is satisfied by an empty trace and returns false otherwise. Thus given a sequence of states $s_1 s_2 \dots s_n$, an LTL formula F written in EAGLE is said to be satisfied by the sequence of states if and only if $value(eval(\dots eval(eval(F, s_1), s_2) \dots s_n))$ is true. The definition of the operators $eval$, $update$ and $value$ forms the calculus of the recursive rule-based framework. We define this calculus next.

3.1 Calculus

The $eval$, $update$ and $value$ operators are defined a priori for all operators. Note that, unlike in general EAGLE where new temporal operators in the form of rules can be defined, in LTL the operators are fixed. So instead of giving a general algorithm to synthesize the definitions of $eval$, $update$ and $value$ for the rules [4], we can synthesize these definitions for the fixed operators of LTL before hand and make them part of our calculus. We do not define the functions on the previous operator, since this operator is eliminated in the the calculus that we present next. The definition of $eval$, $update$ and $value$ on the different operators is given below.

$$\begin{aligned}
eval(jexp, s) &= \text{value of } jexp \text{ in } s \\
eval(F_1 \text{ op } F_2, s) &= eval(F_1, s) \text{ op } eval(F_2, s) \text{ where } op \in \{\wedge, \vee, \rightarrow\} \\
eval(\neg F, s) &= \neg eval(F, s) \\
eval(\bigcirc F, s) &= update(F, s) \\
update(jexp, s) &= jexp \\
update(F_1 \text{ op } F_2, s) &= update(F_1, s) \text{ op } update(F_2, s) \text{ where } op \in \{\wedge, \vee, \rightarrow\} \\
update(\neg F, s) &= \neg update(F, s) \\
update(\bigcirc F, s) &= \bigcirc update(F, s) \\
value(jexp) &= \underline{false} \\
value(F_1 \text{ op } F_2) &= value(F_1) \text{ op } value(F_2) \text{ where } op \in \{\wedge, \vee, \rightarrow\} \\
value(\neg F) &= \neg value(F) \\
value(\bigcirc F) &= \underline{false}
\end{aligned}$$

Note that $eval$ of a formula of the form $\bigcirc F$ on a state s reduces to the $update$ of F on state s . This ensures that if F contains any past time operators then $update$ of F updates them properly. Moreover, $value(\bigcirc F)$ is false as the operator \bigcirc is assumed to have strong interpretation in the logic. The $value$ of a max rule is true and that of a min rule is false.

$$\begin{aligned}
value(R(F_1, \dots, F_n)) &= \underline{true} \text{ if } R \text{ is } \underline{max} \\
value(R(F_1, \dots, F_n)) &= \underline{false} \text{ if } R \text{ is } \underline{min}
\end{aligned}$$

However, the definition of the *eval* and *update* operators for the rules are not generic for all LTL operators. They are synthesized according to the definition of the rules in the specification and made part of the calculus. Consider the Always operator.

$$\underline{\max} \text{ Always}(\underline{\text{Form}} F) = F \wedge \bigcirc \text{Always}(F)$$

For this rule *eval* and *update* are defined as follows.

$$\begin{aligned} \text{eval}(\text{Always}(F), s) &= \text{eval}(F \wedge \bigcirc \text{Always}(F), s) \\ \text{update}(\text{Always}(F), s) &= \text{update}(F \wedge \bigcirc \text{Always}(F), s) \end{aligned}$$

However, the definition of *update* results in infinite recursion. To break the recursion we note that the rule Always does not contain any previous operator, although the argument *F* may contain some. So we simply propagate the *update* to the argument *F*. Thus the new definition of *update* becomes:

$$\text{update}(\text{Always}(F), s) = \text{Always}(\text{update}(F, s))$$

In a similar way we can give the calculus for the other future time LTL operators as follows:

$$\begin{aligned} \text{eval}(\text{Eventually}(F), s) &= \text{eval}(F \vee \bigcirc \text{Eventually}(F), s) \\ \text{update}(\text{Eventually}(F), s) &= \text{Eventually}(\text{update}(F, s)) \end{aligned}$$

$$\begin{aligned} \text{eval}(\text{Until}(F_1, F_2), s) &= \text{eval}(F_2 \vee (F_1 \wedge \bigcirc \text{Until}(F_1, F_2)), s) \\ \text{update}(\text{Until}(F), s) &= \text{Until}(\text{update}(F_1, s), \text{update}(F_2, s)) \end{aligned}$$

$$\begin{aligned} \text{eval}(\text{Unless}(F_1, F_2), s) &= \text{eval}(F_2 \vee (F_1 \wedge \bigcirc \text{Unless}(F_1, F_2)), s) \\ \text{update}(\text{Unless}(F), s) &= \text{Unless}(\text{update}(F_1, s), \text{update}(F_2, s)) \end{aligned}$$

However, the definitions are different for past time LTL operators. These operators defined in the form of rules contain previous operator. In general, if a rule contains a formula *F* guarded by a previous operator on its right hand side then we evaluate *F* at every event and use the result of this evaluation in the next state. Thus, the result of evaluating *F* is required to be stored in some temporary placeholder so that it can be used in the next state. To allocate a placeholder, we introduce, for every formula guarded by a previous operator, an argument in the rule and use these arguments in the definition of *eval* and *update* for this rule. Let us illustrate this with the following example.

$$\underline{\max} \text{ AlwaysInPast}(\underline{\text{Form}} F) = F \wedge \odot \text{AlwaysInPast}(F)$$

For this rule we introduce another auxiliary rule AlwaysInPast' which contains an extra argument corresponding to the formula $\odot(\text{AlwaysInPast}(F))$.

$$\begin{aligned} \text{AlwaysInPast}(\underline{\text{Form}} F) &= \text{AlwaysInPast}'(F, \underline{\text{true}}) \\ \text{eval}(\text{AlwaysInPast}'(F, \text{past}_1), s) &= \text{eval}(F \wedge \text{past}_1, s) \\ \text{update}(\text{AlwaysInPast}'(F, \text{past}_1), s) &= \\ &\quad \text{AlwaysInPast}'(\text{update}(F, s), \text{eval}(\text{AlwaysInPast}'(F, \text{past}_1), s)) \end{aligned}$$

Here, in *eval*, the subformula $\odot(\text{AlwaysInPast}(F))$ guarded by the previous operator is replaced by the argument *past*₁ that contains the evaluation of the subformula in the previous state. In *update* we not only update the argument *F* but also evaluate the subformula $\text{AlwaysInPast}'(F, \text{past}_1)$ and pass it as second argument of $\text{AlwaysInPast}'$. Thus in the next state *past*₁ is bound to $\odot(\text{AlwaysInPast}'(F, \text{past}_1))$. Note that in the definition of $\text{AlwaysInPast}'$ we pass true as the second argument. This is because, AlwaysInPast being defined a maximal operator, its previous value at the beginning of the trace is true.

In a similar way we can give the calculus for the other past time LTL operators as follows:

$$\begin{aligned} \text{Previous}(\underline{\text{Form}} F) &= \text{Previous}'(F, \underline{\text{false}}) \\ \text{eval}(\text{Previous}'(F, \text{past}_1), s) &= \text{eval}(\text{past}_1, s) \\ \text{update}(\text{Previous}'(F, \text{past}_1), s) &= \text{Previous}'(\text{update}(F, s), \text{eval}(\text{Previous}'(F, \text{past}_1), s)) \end{aligned}$$

$$\begin{aligned} \text{EventuallyInPast}(\underline{\text{Form}} F) &= \text{EventuallyInPast}'(F, \underline{\text{false}}) \\ \text{eval}(\text{EventuallyInPast}'(F, \text{past}_1), s) &= \text{eval}(F \vee \text{past}_1, s) \\ \text{update}(\text{EventuallyInPast}'(F, \text{past}_1), s) &= \\ &\text{EventuallyInPast}'(\text{update}(F, s), \text{eval}(\text{EventuallyInPast}'(F, \text{past}_1), s)) \end{aligned}$$

$$\begin{aligned} \text{Since}(\underline{\text{Form}} F_1, \underline{\text{Form}} F_2) &= \text{Since}'(F_1, F_2, \underline{\text{false}}) \\ \text{eval}(\text{Since}'(F_1, F_2, \text{past}_1), s) &= \text{eval}(F_2 \vee (F_1 \wedge \text{past}_1), s) \\ \text{update}(\text{Since}'(F_1, F_2, \text{past}_1), s) &= \\ &\text{Since}'(\text{update}(F_1, s), \text{update}(F_2, s), \text{eval}(\text{Since}'(F_1, F_2, \text{past}_1), s)) \end{aligned}$$

$$\begin{aligned} \text{Zince}(\underline{\text{Form}} F_1, \underline{\text{Form}} F_2) &= \text{Zince}'(F_1, F_2, \underline{\text{true}}) \\ \text{eval}(\text{Zince}'(F_1, F_2, \text{past}_1), s) &= \text{eval}(F_2 \vee (F_1 \wedge \text{past}_1), s) \\ \text{update}(\text{Zince}'(F_1, F_2, \text{past}_1), s) &= \\ &\text{Zince}'(\text{update}(F_1, s), \text{update}(F_2, s), \text{eval}(\text{Zince}'(F_1, F_2, \text{past}_1), s)) \end{aligned}$$

For the sake of completeness of the calculus we explicitly define *value* on the above LTL operators as follows:

$$\begin{array}{ll} \text{value}(\text{Always}(F)) = \underline{\text{true}} & \text{value}(\text{Eventually}(F)) = \underline{\text{false}} \\ \text{value}(\text{Until}(F_1, F_2)) = \underline{\text{false}} & \text{value}(\text{Unless}(F_1, F_2)) = \underline{\text{true}} \\ \text{value}(\text{AlwaysInPast}(F)) = \underline{\text{true}} & \text{value}(\text{EventuallyInPast}(F)) = \underline{\text{false}} \\ \text{value}(\text{Since}(F_1, F_2)) = \underline{\text{false}} & \text{value}(\text{Zince}(F_1, F_2)) = \underline{\text{true}} \end{array}$$

Note that in the above calculus we have got rid of the previous operator by introducing an auxiliary argument or placeholder for every formula guarded by \odot operator. Hence, we cannot use the operator \odot while writing an LTL formula. Instead we use the rule Previous as defined above.

Thus, we translate the rules in the specification to a set of definition of *eval* and *update* operators. Once we have this translation we can easily execute, or in other words, evaluate all the monitors at each state in a trace of a running program.

4 Implementation and Complexity

We have an implementation for the monitoring framework for EAGLE in Java. The implemented system works in two phases. First, it compiles the specification file to *synthesize* a set of Java classes; a class is generated for each rule. Second, the Java class files are compiled into Java bytecode and then the monitoring engine dynamically loads the Java classes for rules at monitoring time and monitors a trace.

However, for the purpose of LTL monitoring we do not have to synthesize the Java classes as the set of rules are fixed. Rather, we hardwire the whole algorithm in the implementation.

In order to make the implementation efficient we use the decision procedure of Hsiang [13]. The procedure reduces a tautological formula to the constant true, a false formula to the constant false, and all other formulas to canonical forms which are exclusive disjunction (\oplus) of conjunctions. The procedure is given below using equations that are shown to be Church-Rosser and terminating modulo associativity and commutativity.

$$\begin{array}{ll}
 \text{true} \wedge \phi = \phi & \text{false} \wedge \phi = \text{false} \\
 \phi \wedge \phi = \phi & \text{false} \oplus \phi = \phi \\
 \phi \oplus \phi = \text{false} & \neg \phi = \text{true} \oplus \phi \\
 \phi_1 \wedge (\phi_2 \oplus \phi_3) = (\phi_1 \wedge \phi_2) \oplus (\phi_1 \wedge \phi_3) & \phi_1 \vee \phi_2 = (\phi_1 \wedge \phi_2) \oplus \phi_1 \oplus \phi_2 \\
 \phi_1 \rightarrow \phi_2 = \text{true} \oplus \phi_1 \oplus (\phi_1 \wedge \phi_2) & \phi_1 \equiv \phi_2 = \text{true} \oplus \phi_1 \oplus \phi_2
 \end{array}$$

In particular the equations $\phi \wedge \phi = \phi$ and $\phi \oplus \phi = \text{false}$ ensures that, at the time of monitoring, we do not expand the formula beyond bound. The bound is given by the following theorem:

Theorem 1. *The size of the formula at any stage of monitoring is bounded by $O(\text{size}(\phi) \cdot 2^{\text{size}(\phi)})$, where ϕ is the initial LTL formula for which we started monitoring.*

Proof. The above equations, when regarded as simplification rules, keeps any LTL formula in a canonical form, which is an exclusive disjunction of conjunctions, where conjuncts have temporal operators at top. Moreover, after a series of applications of *eval* on the states s_1, s_2, \dots, s_n , the conjuncts in the normal form $\text{eval}(\dots \text{eval}(\text{eval}(\phi, s_1), s_2) \dots, s_n)$ are subterms of the initial formula ϕ , each having a temporal operator at its top. Since there are at most $\text{size}(\phi)$ such subformulas, it follows that there are at most $2^{\text{size}(\phi)}$ possibilities to combine them in a conjunction. The space requirement for each conjunct is $\text{size}(\phi)$. Therefore, one needs space $O(\text{size}(\phi) \cdot 2^{\text{size}(\phi)})$ to store any exclusive disjunction of such conjunctions. \square

The implementation contains a strategy for the application of these equations that ensures that the time complexity of each step in monitoring is $O(\text{size}^2(\phi) \cdot 2^{\text{size}(\phi)})$. We next describe the strategy briefly. Since, our LTL formulas are exclusive disjunction of conjunctions we can treat them as a tree of depth two: the root node at depth 0 representing the \oplus operator, the children of the root at depth 1 representing the \wedge operators, and the leaf nodes at depth 2 representing the temporal operators and the Java expressions. For example, figure 1 shows the tree representation of the formula $p \rightarrow \diamond(q \mathcal{U} r)$, whose canonical form is $\text{true} \oplus p \oplus (p \wedge \diamond(q \mathcal{U} r))$.

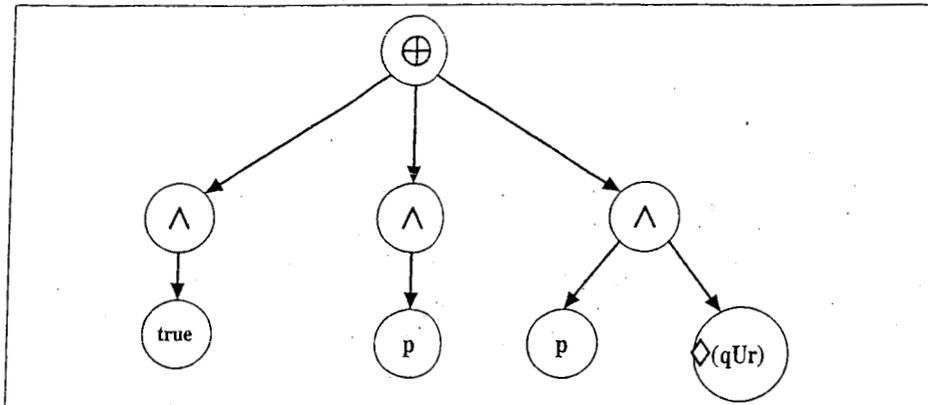


Fig. 1. Tree representation of $p \rightarrow \diamond(q \cup r)$

When we apply *eval* on a formula and a state the *eval* function is applied in depth-first fashion on this tree and we build up the resultant formula in a bottom-up fashion. At the leaves the application of *eval* results either in the evaluation of a Java expression or the evaluation of a rule. The evaluation of a Java expression returns either true or false. We assume that this evaluation takes unit time. On the other hand, the evaluation of a rule may result in another formula in canonical form. The formula at any internal node is then evaluated by taking the conjunction (or exclusive disjunction) of the formulas of the children nodes as they get evaluated. The following gives the pseudocode for the strategy:

```

Form eval(F,s)
begin
  Form F';
  if F is conjunction of subformulas then
    F' = true;
    for each subformula Fsub of F do
      F' = F'  $\wedge$  eval(Fsub,s);
    endfor
  else if F is exclusive disjunction of subformulas then
    F' = false;
    for each subformula Fsub of F do
      F' = F'  $\oplus$  eval(Fsub,s);
    endfor
  else if F is a rule or expression then
  endif
  return F';
endsub

```

Note that the application of conjunction on two formulas in canonical form requires the application of the distributive equation $\phi_1 \wedge (\phi_2 \oplus \phi_3) = (\phi_1 \wedge \phi_2) \oplus (\phi_1 \wedge \phi_3)$ and possibly other equations.

At any stage of this algorithm there are three formulas that are active: the original formula F on which $eval$ is applied, the formula F' , and the result of the evaluation of the subformula $Fsub$. So, by theorem 1 we can say that the space complexity of this algorithm is $O(size(\phi).2^{size(\phi)})$. Moreover, the algorithm traverses the formula once at each node it can possibly spend $O(size(\phi).2^{size(\phi)})$ time to do the conjunction and exclusive disjunction. Hence the time complexity of the algorithm is $O(size(\phi).2^{size(\phi)})$, $O(size(\phi).2^{size(\phi)})$ or $O(size^2(\phi).2^{size(\phi)})$. These two bounds are given as the following theorem.

Theorem 2. *At any stage of monitoring the space and time complexity of the evaluation of the monitored LTL formula on the current state is $O(size(\phi).2^{size(\phi)})$ and $O(size^2(\phi).2^{size(\phi)})$ respectively.*

EAGLE has been applied to test a planetary rover controller in a collaborative effort with other colleagues, see [2] for an earlier similar experiment using a simpler logic. The rover controller, written in 35,000 lines of C++, executes action plans. The testing environment, consists of a test-case generator, automatically generating input plans for the controller. Additionally, for each input plan a set of temporal formulas is generated that the plan execution should satisfy. The controller is executed on the generated plans and the implementation of EAGLE is used to monitor that execution traces satisfy the formulas. A previously unknown error was detected in the first run, demonstrating that a certain task did not recognize the too early termination of some other task.

5 Conclusion and Future Work

We have presented a representation of linear temporal logic with both past and future temporal operators in EAGLE. We have shown how the generalized monitoring algorithm for EAGLE becomes simple and elegant for this particular case. We have bounded the space and time complexity of this specialized algorithm and thus showed that general LTL monitoring is efficient if we use EAGLE framework. Initial experiments have been successful. Future work includes: optimizing the current implementation; investigating other efficient subsets of EAGLE and associating actions with formulas.

References

1. *1st, 2nd and 3rd CAV Workshops on Runtime Verification (RV'01 - RV'03)*, volume 55(2), 70(4), 89(2) of *ENTCS*. Elsevier Science: 2001, 2002, 2003.
2. C. Artho, D. Drusinsky, A. Goldberg, K. Havelund, M. Lowry, C. Pasareanu, G. Roşu, and W. Visser. Experiments with Test Case Generation and Runtime Analysis. In E. Börger, A. Gargantini, and E. Riccobene, editors, *Abstract State Machines (ASM'03)*, LNCS, pages 87–107. Springer, March 2003.
3. H. Barringer, M. Fisher, D. Gabbay, G. Gough, and R. Owens. METATEM: An introduction. *Formal Aspects of Computing*, 7(5):533–549, 1995.

4. H. Barringer, A. Goldberg, K. Havelund, and K. Sen. Rule-based runtime verification. Manuscript, August 2003.
5. D. Drusinsky. The Temporal Rover and the ATG Rover. In K. Havelund, J. Penix, and W. Visser, editors, *SPIN Model Checking and Software Verification*, volume 1885 of *LNCS*, pages 323–330. Springer, 2000.
6. D. Drusinsky. Monitoring Temporal Rules Combined with Time Series. In *CAV'03*, volume 2725 of *LNCS*, pages 114–118. Springer-Verlag, 2003.
7. B. Finkbeiner, S. Sankaranarayanan, and H. Sipma. Collecting Statistics over Runtime Executions. In *Proceedings of Runtime Verification (RV'02)* [1], pages 36–55.
8. B. Finkbeiner and H. Sipma. Checking Finite Traces using Alternating Automata. In *Proceedings of Runtime Verification (RV'01)* [1], pages 44–60.
9. D. Gabbay. The Declarative Past and Imperative Future: Executable Temporal Logic for Interactive Systems. In *Proceedings of the 1st Conference on Temporal Logic in Specification, Altrincham, April 1987*, volume 398 of *LNCS*, pages 409–448, 1989.
10. D. Giannakopoulou and K. Havelund. Automata-Based Verification of Temporal Properties on Running Programs. In *Proceedings, International Conference on Automated Software Engineering (ASE'01)*, pages 412–416. ENTCS, 2001. Coronado Island, California.
11. K. Havelund and G. Roşu. Monitoring Programs using Rewriting. In *Proceedings, International Conference on Automated Software Engineering (ASE'01)*, pages 135–143. Institute of Electrical and Electronics Engineers, 2001. Coronado Island, California.
12. K. Havelund and G. Roşu. Synthesizing Monitors for Safety Properties. In *Tools and Algorithms for Construction and Analysis of Systems (TACAS'02)*, volume 2280 of *Lecture Notes in Computer Science*, pages 342–356. Springer, 2002.
13. Jieh Hsiang. Refutational Theorem Proving using Term Rewriting Systems. *Artificial Intelligence*, 25:255–300, 1985.
14. D. Kortenkamp, T. Milam, R. Simmons, and J. Fernandez. Collecting and Analyzing Data from Distributed Control Programs. In *Proceedings of RV'01* [1], pages 133–151.
15. K. Jelling Kristoffersen, C. Pedersen, and H. R. Andersen. Runtime Verification of Timed LTL using Disjunctive Normalized Equation Systems. In *Proceedings of Runtime Verification (RV'03)* [1], pages 146–161.
16. I. Lee, S. Kannan, M. Kim, O. Sokolsky, and M. Viswanathan. Runtime Assurance Based on Formal Specifications. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications*, 1999.
17. A. Pnueli. The Temporal Logic of Programs. In *Proceedings of the 18th IEEE Symposium on Foundations of Computer Science*, pages 46–77, 1977.
18. K. Sen and G. Roşu. Generating Optimal Monitors for Extended Regular Expressions. In *Proceedings of the 3rd Workshop on Runtime Verification (RV'03)* [1], pages 162–181.